

Architecture, Design, Implementation

Amnon H. Eden

*Center for Inquiry, Amherst, NY, and
Department of Computer Science,
University of Essex, United Kingdom*

eden@acm.org

Rick Kazman

*Software Engineering Institute, Pittsburgh, PA
and University of Hawaii, Honolulu, HI*

kazman@sei.cmu.edu

Abstract

The terms architecture, design, and implementation are typically used informally in partitioning software specifications into three coarse strata of abstraction. Yet these strata are not well-defined in either research or practice, causing miscommunication and needless debate.

To remedy this problem we formalize the Intension and the Locality criteria, which imply that the distinction between architecture, design, and implementation is qualitative and not merely quantitative. We demonstrate that architectural styles are intensional and non-local; that design patterns are intensional and local; and that implementations are extensional and local.

*If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is.
— J. H. von Neumann*

1. Introduction

In their seminal article, Perry and Wolf [24] developed “an intuition about software architecture through analogies to existing disciplines.” Building on this, Shaw and Garlan [31] suggest that “software architecture involves the description of elements from which systems are built.” A considerable body of work, stemming back to DeRemer and Kron’s *module interconnection languages* (MIL) [7], focuses on the specification, construction, and analysis of large software systems defined by these terms (e.g., [26], [21], [15]). For example, *architecture description languages* (ADL) combine a formal specification language with tools supporting the construction and analysis of software systems from such specifications.

Seeking to separate *architectural design* from other design activities, definitions of *software architecture* stress the following:

- “*Architecture* is concerned with the selection of architectural elements, their interaction, and the constraints on those elements and their interactions... *Design* is concerned with the modularization and detailed interfaces of the design elements, their algorithms and pro-

cedures, and the data types needed to support the architecture and to satisfy the requirements.” [24]

- Software Architecture is “concerned with issues ... beyond the algorithms and data structures of the computation.” [16]
- “Architecture ... is specifically not about ... details of implementations (e.g., algorithms and data structures.) ... Architectural design involves a richer collection of abstractions than is typically provided by OOD.” [23]
- “Architecture \neq Design? ... Design is an activity. Architecture, or architectural design, is design at a higher level of abstraction.” [18]
- Architecture focuses on the externally visible properties of software “components.” [2]

In suggesting typical “architectures” and “architectural styles”, existing definitions consist of examples and offer anecdotes rather than provide unambiguous, clear notions.

In practice, the terms “architecture”, “design” and “implementation” appear to connote varying degrees of abstraction in the continuum between complete details (“implementation”), few details (“design”), and the highest form of abstraction (“architecture”). But the amount of detail alone is insufficient to characterize the differences, because architecture and design documents often contain information that is not explicit in the implementation (e.g., design constraints, standards, performance goals) and therefore they cannot result from mere omission of detail. Thus, we would expect a distinction to be qualitative and not merely quantitative. A clear distinction has remained elusive and this lack of distinction is the cause of much muddy thinking, imprecise communication, and wasted, overlapping effort.

As a result, *architecture* is often used as a mere synonym for *design*. For example, the “Siemens” catalogue [4] defines “architectural patterns” that are in par with “design patterns” defined by the “Gang of Four” [14].

Confusion also stems from the use of the same specification language for both architectural and design specifications. For example, the *Software Engineering Institute* (SEI) classifies UML [3] as an architectural description language [30], and it has become the industry de facto standard ADL, although UML was specifically designed

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAY 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE Architecture, Design, Implementation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,Software Engineering Institute,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The terms architecture, design, and implementation are typically used informally in partitioning software specifications into three coarse strata of abstraction. Yet these strata are not well-defined in either research or practice causing miscommunication and needless debate. To remedy this problem we formalize the Intension and the Locality criteria, which imply that the distinction between architecture, design, and implementation is qualitative and not merely quantitative. We demonstrate that architectural styles are intensional and non-local; that design patterns are intensional and local; and that implementations are extensional and local.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

to manifest detailed design decisions (and this is its most common use).

Confusion also exists with respect to the artifacts of design and implementation. UML *class diagrams* [3], for instance, are a prototypical artifact of the design phase. Nonetheless, class diagrams may accumulate enough detail to allow code generation of very detailed programs, an approach that is promoted by CASE tools such as *Rational Rose*® [28] and *System Architect*® [25]. Using the same specification language further blurs the distinction between artifacts of the design (class diagrams) from the implementation (source code.)

Intended contribution. Why are we interested in such distinctions? Naturally, a well-defined language improves our understanding of the subject matter. With time, terms that are used interchangeably lose their meaning and end up as mere platitudes, resulting inevitably in ambiguous descriptions given by developers, and significant effort is wasted in discussions of the form “by design I mean... and by architecture I mean...” The formal ontology we provide can serve as the ultimate reference point for these discussions.

The contribution of this paper is to provide insight on the largely informal dialectic by appealing to both intuition and to formal ontology. By putting these terms on a solid footing not only do we disambiguate the progressively murky discourse in “architectural specifications” but provide a foundation for formal reasoning and analysis, as well as a firm foundation for informal “chalk-talk” discussions. Finally, tools supporting design and architectural specifications, where intuitive perceptions are insufficient, will benefit by accurately defining this distinction.

Many of the required definitions were rendered informal and the proofs were omitted in this version of our work. The interested reader can find the complete definitions in [12], and additional details on the *design models* formalism can be found in [9]. Instead, we argue our points with informal illustrations and discussions.

1.1 The Intension/Locality Thesis

To elucidate the relationship between *architecture*, *design*, and *implementation*, we distinguish at least two separate interpretations for *abstraction* in our context:

Intensional (vs. *extensional*) specifications are “abstract” in the sense that they can be formally characterized by the use of logic variables that range over an unbounded domain;

Non-local (vs. *local*) specifications are “abstract” in the sense that they pervade *all* parts of the system (as opposed to being limited to some part thereof).

Both of these interpretations contribute to the distinction among *architecture*, *design*, and *implementation*. This combination of these interpretations leads us to the *intension/locality thesis*:

- (i) Architectural specifications are *intensional* and *non-local*;
- (ii) Design specifications are *intensional* but *local*; and
- (iii) Implementation specifications are both *extensional* and *local*.

The intension/locality thesis is summarized, for easier reference, in Table 1.

Table 1. The Intension/Locality Thesis

Architecture	<i>Intensional</i>	<i>Non-local</i>
Design	<i>Intensional</i>	<i>Local</i>
Implementation	<i>Extensional</i>	<i>Local</i>

1.2 Structure of This Paper

The intension/locality thesis can be understood correctly only in the context of the ontology provided below. In Section 2 we define *design models*, which are crucial to the remainder of our discussion. Design models are abstractions which allow a formal “meaning” assigned to programs, also called “denotation”. This formalism allows us to determine whether a specification is “satisfied” by a program.

In Section 3, we formally define the Intension criterion and the Locality criterion. We distinguish our interpretation for “intensionality” from the accepted usage, as we define it in terms of *design models*.

Sections 4, 5, and 6, provide case studies in applying the Intension and Locality criteria using our formal ontology. In Section 4 we demonstrate that implementations in any programming language, including generics and C++ templates are *extensional* and *local*. In Section 5 we show that design patterns, such as the Factory Method, and design specifications, such as the *Enterprise JavaBeans*™ and Java™ *Swing*’s MVC, are *intensional* and *local*. In Section 6 we demonstrate that architectural styles such as Pipes and Filters and Layered Architecture are *intensional* and *non-local*, and so is the Law of Demeter.

In Section 7, we discuss some of the ramifications of our criteria. The discussion of UML class diagrams in Section 7.2 reveals that class diagrams have a separate place in the hierarchy of abstractions we describe. Section 8 summarizes the contributions of this paper.

2. Setting the Scene

In this section, we illustrate the formal ontology that underlies our discussion. This ontology is based on giving an abstract “meaning” to programs using *design models*.

2.1 Design Models

Turing and *random-access machines* provide robust computational models that are primarily suitable for reasoning about algorithms. Other computational models and formalisms (e.g., Petri nets, statecharts, and temporal logic) facilitate reasoning about certain behavioral properties.

The discussion in architectural and design specifications, however, involves reasoning on constructs such as *classes*, *methods*, and *function calls*. Most other formalisms incorporate too much implementation detail and do not allow a discussion in the appropriate level of abstraction. As we seek to establish the relation between architectural or design specifications and implementations, we base our discussion on a different formalism, one which abstracts programs to a more convenient representation.

Eden and Hirshfeld [11] demonstrate how to model source code as *design models*, which are first order, finite

structures in mathematical logic [1]. Informally, a **design model** m [11] consists of a set of *atoms* and a set of *relations* among those atoms.

Table 2 depicts a detailed example of a trivial Java™ program and a design model that represents it. As this example demonstrates, an object-oriented program is abstracted to a collection of *classes* and *methods* (also *routines* or *function members*) and their relations. *Atoms* represent *classes* and *methods* that were defined in the program, such as the class `Decorator` and the method `Decorator.Draw`. *Relations* represent their correlations, such as

$$\text{Inherit}(\text{BorderDecorator}, \text{Decorator}) \quad (1)$$

Note that *design models* are abstractions which were made to reflect certain structural aspects of computer programs that are relevant to the discussion in software design theory. Obviously, this representation limits the type of reasoning we may perform to properties that are relevant to the discussion in classes, methods, and their interdependencies, as opposed to the discussion in dynamic properties such as *fairness* and complexity.

2.2 Specifications and Instances

In this subsection, we discuss specifications and programs, and illustrate how the two correlate. We make some reasonable assumptions on the languages used to write specifications.

Let us designate $SPEC$ as the set of formal languages of any order [1]. Let $SPEC^*$ designate the set of all expressions made in some language in $SPEC$. A **specification** is an element of $SPEC^*$.

$SPEC$ includes familiar specification languages such as \mathbb{Z} [32], as demonstrated in formulas (5.1) and (5.2), and LePUS [8], as demonstrated in formula (4). $SPEC$ also includes programming languages such as Eiffel, C++, and Java™. Naturally, $SPEC$ is not restricted to known programming or specification languages.

A specification is only useful if we can determine whether it is “satisfied” or not. Having chosen *design models* as our semantics we wish to ask: Does this program satisfy our specification? To answer this question, consider for example the following trivial design specification:

$$\text{Inherit}(x, y) \quad (2)$$

Expression (2) contains two free variables: x , y . We say that it can be *satisfied* by any pair of atoms that belong to the relation *Inherit*. For example, from expression (1) we conclude that the pair $\langle \text{BorderDecorator}, \text{Decorator} \rangle$ satisfies expression (2).

Table 2. A Java™ program and its denotation (adapted from [9].)

<pre> abstract class Decorator { public void Draw(); } class BorderDecorator extends Decorator { public void Draw() { Decorator.Draw(); } private int BorderWidth; } </pre>
<p>The design model of this program consists of the following:</p> <p><u>Atoms:</u></p> <p>“class” atoms: {Decorator, BorderDecorator, int, void}, “Method” atoms: {BorderDecorator.Draw, Decorator.Draw}</p> <p><u>Relations:</u></p> <p><i>Abstract</i>(Decorator)</p> <p><i>Member</i>(Decorator.Draw, Decorator)</p> <p><i>Member</i>(BorderDecorator.Draw, BorderDecorator)</p> <p><i>Inherit</i>(BorderDecorator, Decorator)</p> <p><i>Reference</i>(BorderDecorator, int)</p> <p><i>Invoke</i>(BorderDecorator.Draw, Decorator.Draw)</p> <p><i>ReturnType</i>(Decorator.Draw, void)</p> <p><i>ReturnType</i>(BorderDecorator.Draw, void)</p>

Thus, we can say that the pair of atoms $\langle \text{BorderDecorator}, \text{Decorator} \rangle$ is an **instance** of (2), and that the design model depicted in Table 2 **instantiates** expression (2). A more formal definition of *instance* appears in [12] and in [9].

From this example, it should be clear that an *instance* is not the same as a “program”. Depending on the specification, a program may contain zero, one, or any number of instances. The following subsection formalizes the notion of a program.

2.3 Programs and their Denotation

We expect a “program” to be a specification that is associated with only one *design model*. The association between “real” programs and design models is provided by the *denotation function*. Loosely speaking, a **denotation function** \mathbf{D} has these properties:

- Its domain, denoted \mathcal{P}_D , is a subset of $SPEC^*$;
- \mathbf{D} associates each element φ with *exactly one* design model (its *denotation*) which *instantiates* φ .

Typically, the domain of \mathbf{D} contains expressions in programming languages, such as C++ and Eiffel. More formally, a **program** is an element in \mathcal{P}_D . Every design model typically denotes infinitely many programs. Figure 1 illustrates the denotation associated between certain specifications (“programs”) and design models. Table 2 illustrates the denotation of a simple Java™ program and Table 3 the denotation of a C++ program.

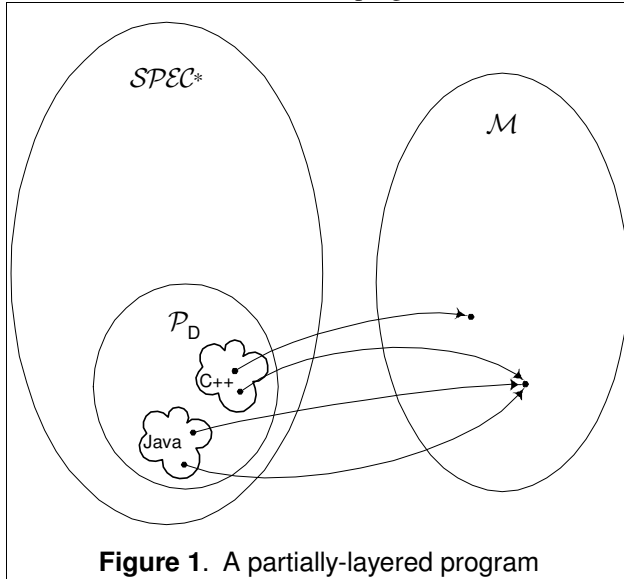


Figure 1. A partially-layered program

Observe that the set of programs \mathcal{P}_D is only a *small subset* of the expressions in $SPEC^*$, and that there are many expressions in $SPEC^*$ that are instantiated by more than one design model

We now have a well-defined notion of programs and specifications. In combination with the definition of an “instance” of a specification, we can conclusively determine whether a program satisfies a given specification:

Definition I. We say that a program π *satisfies* specification φ iff the design model of π instantiates φ .

This means that, for example, the Java™ program depicted in Table 2 *satisfies* expression (2).

In the following sections, we will set apart architecture, design, and implementation specifications based on observing properties of the groups of programs that satisfy each specification.

3. The Intension/Locality Criteria

We will now define the concepts of *intension* and *locality*, which will be applied, both formally and informally, to selected specifications in the following sections.

3.1 The Intension Criterion

Perry and Wolf [24] have established that architectural specifications must be made in intensional terms. Speaking of the desired properties of an ideal specification language for software architecture they write: “We want a means of supporting a ‘principle of least constraint’ to be able to express only those constraints in the architecture that are necessary at the architectural level of the system description”. It constrains only what it needs to, in terms of properties imposed over free variables.

Traditionally, *intensional specifications* define a concept via a list of constraints. For example, mathematical concepts are usually defined intensionally. For instance, “A *prime number* is a number that divides only by itself and by the number 1”. In contrast, the North Atlantic Treaty defines the set of NATO members extensionally, namely, by itemizing its members: United States, United Kingdom, Norway, and so forth.

The notion of intensionality defined below diverges slightly from the philosophical concept. We say that a specification is intensional if and only if it has an unbounded number of instances:

Definition II. We say that a specification is *intensional* iff there are infinitely-many possible instances thereof. Conversely, all other expressions are *extensional*.

It immediately follows (as shown in [12]) that intensional specifications are also satisfied by infinitely many design models and by infinitely many programs.

3.2 The Locality Criterion

Monroe et. al [23] argue that “Architectural designs are typically concerned with the entire system.” Similarly, we observe that an architectural style, which pervades a system [16], manifests a property that is shared across modules of the system. This intuition motivates the Locality criterion: What distinguishes architectural from design specifications is that *architectural specifications must be met by every extension of the program*.

As a simple example, consider the rule of a “universal base class”. Although the language does not require it, several C++ class libraries (e.g., NIHCL and Microsoft’s MFC) are designed by this rule. Formally, this property can be expressed as follows:

$$\forall c \bullet \text{Class}(c) \Rightarrow \text{Inherit}^*(c, \text{Object}) \quad (3)$$

(Where Inherit^* is the transitive closure of the binary relation Inherit .) The intension/locality thesis argues that Universal Base Class is architectural because (a) it is intensional and (b) it pervades *all parts* of the system, i.e., *every* class must be bound to `Object`.

Subsumption. The formalization of the locality criterion requires the notion of *subsumption*. Informally, we say that structure \mathbf{n} *subsumes* structure \mathbf{m} if \mathbf{m} is a “sub-model” of \mathbf{n} , or that \mathbf{n} is an “extension” to \mathbf{m} .

Definition III. We say that a specification φ is **local** iff the following condition holds:

If φ is satisfied in some design model \mathbf{m} then every design model that subsumes \mathbf{m} also satisfies φ .

Essentially, Definition III states that an expression is *local* if it can be satisfied in “some corner” of our program without this being affected in how the rest of the program is like.

In the following sections, we apply the Intension and Locality criteria to selected specifications to illustrate the difference between programs, design specifications, and architectural specifications.

4. Implementations

It immediately follows from the Intension criterion (as shown in [12]) that intensional specifications are also satisfied by infinitely many design models and by infinitely many programs. This leads us to prove part (iii) of the intension/locality thesis:

The lemma of “extensions”: *Programs are extensional specifications.*

Following the ‘principle of least constraint’, we expect architectural specifications to have an unbounded number

of instances, namely, to be “intensional”. The same applies to design patterns. But what about other forms of specifications? Can programs be intensional?

Prima facie, it appears that some programming specifications (such as C++ templates and Eiffel generics) might also be intensional. This is not true in the context of *design models*: As demonstrated in Corollary 1, specifications in any programming language, including generics and interpreted code are, under the assumptions provided above, purely extensional:

Corollary 1. *C++ templates are extensional.*

To illustrate this, consider the design model of a C++ program with templates, such as shown in Table 3.

Table 3. A C++ program and its denotation

<pre> template <class C> class Stack { /* ... */ } int main() { Stack<int> si; return 0; } </pre>
<p>This program is interpreted by only one design model, which consists of the following:</p> <p><u>Atoms:</u> “Class” atoms: {Stack, si} “Method” atoms: {main}</p> <p><u>Relations:</u> Generic(Stack) Instantiate(Stack, si, int) Return(main, int)</p>

Corollary 1 demonstrates that although generics may be viewed as intensional with respect to other semantic frameworks, e.g., because they can be used to define other concrete constructs, the ontology we have provided assigns then with only one “interpretation”. The reason is that the formal semantics we chose for the representation of programs are *design models*, which are more abstract than the machine code generated from compilation and from other formal frameworks.

To recap, the formal framework provided in Section 2 guarantees that expressions in all conventional programming languages are extensional.

5. Design Specifications

By part (ii) of the intension/locality thesis, design specifications should be local and intensional. Since design specifications are, in practice, defined informally, we begin section with exploring the intuition behind our thesis. For the purpose of the formal analysis which follows, however, we make use of formal specifications and apply them to widely recognized designs.

5.1 Design Patterns

In this subsection, we focus on an example drawn from the published patterns literature. This allows us to test our ideas on some of the most widely used design specifications.

Coplien and Schmidt [5] argue that “design patterns capture the static and dynamic structures of solutions that occur repeatedly when producing applications in a particular context”. Stripped from the context of a particular application, design patterns represent *categories* of solutions, each pattern has an unbounded number of implementations (as implied by the very choice of the name “pattern”). Thus, they are expected to be *intensional*.

Design patterns are commonly perceived as “less abstract” than architectural specifications. For example, they are commonly referred to as “microarchitectures” [29], that is, as if they were like architectures that only apply to a limited module. Using our terminology, we thus expect them to be *local*.

Consider, for example, the Factory Method design pattern [14]. Essentially, the pattern’s solution offers three sets of participants:

1. A set of *product* classes
2. A set of *factory* classes
3. A set of *factory* methods

The collaborations between these participants are constrained as follows:

4. All *factory methods* share the same signature (thereby allowing for dynamic binding), and each is defined in a different *factory* class.
5. Each *factory method* produces instances of exactly one *products* class.

Figure 2 illustrates the general notion of the pattern. Observe that the set of $\langle \text{factory-}i, \text{factory-method-}i, \text{product-}i \rangle$ triplets is unbounded, because the number of possible *factory* and *product* classes is not bounded.

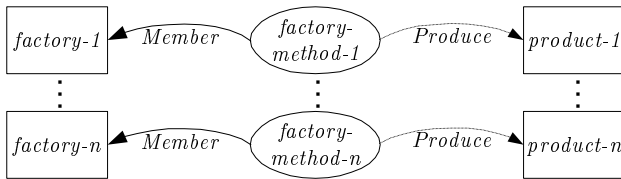


Figure 2. The general structure of the “Factory Method” pattern

For the discussion in design patterns we use LePUS, a formal specification language for object-oriented design, which is defined in detail in [8]. The five statements in the above listed definition of the Factory Method are formally expressed by expressions (4.1) to (4.5) as follows:

$$\text{Products} : \mathbf{P}(\mathbb{C}) \quad (4.1)$$

$$\text{Factories} : \mathbf{P}(\mathbb{C}) \quad (4.2)$$

$$\text{FactoryMethods} : \mathbf{P}(\mathbb{F}) \quad (4.3)$$

$$\text{Clan}(\text{FactoryMethods}, \text{Factories}) \quad (4.4)$$

$$\text{Produce}^{\rightarrow}(\text{FactoryMethods}, \text{Products}) \quad (4.5)$$

Expressions (4.1) through (4.3) declare two sets of *classes* and one set of *methods*. Expression (4.4) indicates that the pair $\langle \text{FactoryMethods}, \text{Factories} \rangle$ satisfies the predicate *Clan*, meaning that –

- All “factory methods” share the same signature (i.e., they share the same dispatch table);
- The relation *MemberOf* is a bijection (i.e., *one-to-one* and *onto* function) between the sets *FactoryMethods* and *Factories*.

Finally, expression (4.5) indicates that the ground relation *Produce* is a bijective function between the set *FactoryMethods* and the set *Products*.

Corollary 2. “Factory method” is *intensional* and *local*.

The composite expression (4.1) to (4.5) is evidently *intensional*, since each one of the free variables *Products* and *Factories* (*FactoryMethods*) can be instantiated by any number of classes (methods). To show that it is *local*, observe that if a design model \mathbf{m} satisfies it then it incorporates an *instance* of the pattern. It is then easy to show that any proper “extension” to \mathbf{m} (i.e., any design model that *subsumes* it) also contains the same instance of the Factory Method, namely, the same set of atoms and relations that satisfied the pattern in \mathbf{m} .

The same line of reasoning can be applied to the specifications of most of the design patterns from the Gamma et. al [14] catalogue, such as the specifications in [8].

5.2 Other Design Specifications

The place of other design specifications within the intension/locality classification may be less obvious than that of design patterns. Thus, we have carried out our analysis on the formal rendering of two additional design specifications: MVC (Model-View-Controller) “usage pattern” in Java™ *Swing* class library, and of Enterprise JavaBeans™.

In lack of space, we cannot quote here the formal specifications but only the results of our analysis. The interested reader may find both specifications in [10], and the complete proof for this conclusion in [12]. We can report that our analysis confirms that, as predicted by the intension/locality thesis, both specifications fall under the “design” category, namely, they are *intensional* and *local*.

6. Architectural Specifications

In this section, we demonstrate that, as predicted by the intension/locality thesis, two classic architectural styles are both intensional and non-local. We also demonstrate that the Law of Demeter is architectural.

6.1 Layered Architecture

Garlan and Shaw [16] describe the *layered architecture* such that “An element of layer k may depend only on elements of layers $1, \dots, k$.” We may formalize this description in \mathbb{Z} as follows:

$$\forall e \exists! k \in \{\mathbb{N}\} \bullet \text{Layer}(e) = k \quad (5.1)$$

(i.e., each element is defined in exactly one layer), and

$$\forall x, y \bullet \quad (5.2)$$

$$\text{Depends}(x, y) \Rightarrow \text{Layer}(x) \geq \text{Layer}(y)$$

(i.e., the definition of each element may “depend” only on the definition of elements of same layer or of lower layers.)

Corollary 3. “Layered architecture” is intensional and non-local.

It is obvious that an unbounded number of programs can satisfy the conjunction of formulas (5.1) and (5.2), hence it is intensional. To prove that it is non-local we show that, given any design model that satisfies this style, the same design model can be extended with a new element in the lower “layer” such that this element depends on a higher layer, thereby violating the specification.

We conclude that we may selectively apply a non-local specification only to certain parts of a program. But what does it mean? Does it compromise our results?

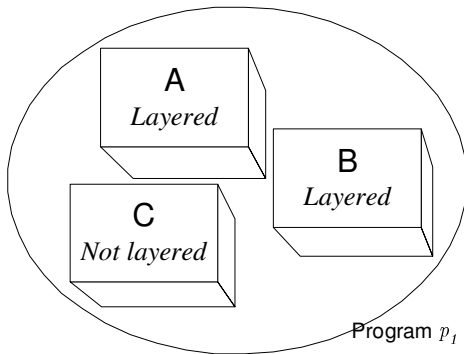


Figure 3. A partially-layered program. Only some parts of program p_1 satisfy Layered Architecture, but p_1 as a “program” does not satisfy the architectural style.

Discussion. Is it possible for a non-local specification, such as layered architecture, to be applied only by one part of a program? Obviously. That simply means that parts of this program satisfy the non-local rule, while other parts violate it. In fact, this property is exactly what makes the layered architecture non-local: We say that it is non-local *because* it may be violated anywhere. Figure 3 and the discussion which follows it illustrate this argument using the Layered Architecture style. Does it mean that non-locality is meaningless? Not at all. Non-locality is a property of a *specification*, i.e., of an expression; and by our thesis, architectural specifications are always non-local. However, an architect may choose to apply a specification only to some parts of the system, and violate it in others.

To summarize, it is not meaningless or contradictory to state that architectural specifications can be applied selectively such that they are violated by some parts of a specific program. In such a case, we say that the style no longer characterizes this program (as a whole.) Formally:

Corollary 4. If a specification φ is (deliberately) violated in module m of program π , then either one of the following is true:

- φ is not satisfied by π , or
- m is not considered part of π (i.e., it belongs to a separate program.)

In the example of Layered Architecture (Figure 3), a module that does “layer bridging” (i.e., violates the layering principle) should not be considered as part of the layered program; instead, we perceive it a separate program.

While this conclusion may seem counter-intuitive at first, it is actually a powerful view on exceptions to architectural constraints. A module that does layer bridging requires different reasoning and management than the rest of the layered system. It not only should, but also *must*, be treated as an exception, or else the power of the layering is compromised. Exceptions to an architectural style should have attention called to them and be made the focus of intense analysis. Our reasoning provides a sound basis for saying when a portion of a program is an exception to an architectural style.

6.2 Pipes and Filters

According to Garlan and Shaw [16], “In a pipes and filter style each component has a set of inputs and a set of outputs. A component reads streams of data on its inputs and produces streams of data on its outputs.” Dean and Cordy [6] present a visual formalism defined as a context-free grammar, and formulate the pipes and filter style as depicted in Figure 4. Their specification is explained in Figure 5 in more intuitive terms.

More generally, a “program” in STSA (*) is represented as a typed, directed multigraph. An *architectural style* is defined as a context-free language. Thus, an expression in STSA defines a collection of graphs in a visual notation, such as Figure 4. According to [6], a “program” satisfies the pipes and filters architecture if it “parses” Figure 4. Figure 5 illustrates the kind of programs that parse the grammar defined in Figure 4.

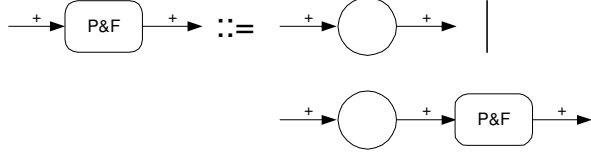


Figure 4. Pipes and Filters (adapted from [6]): Circles represent tasks, arrows represent streams. The plus sign is the BNF symbol for “one or more.”

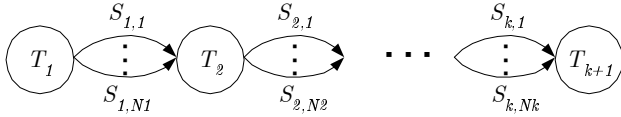


Figure 5. The general structure of programs that parse Figure 4.

Before we can reason about Figure 4, we must establish which design models satisfy an STSA diagram. For the purpose of our discussion in this section, it suffices to restrict ourselves to multigraphs that contain “tasks” and “streams”. More specifically, let G designate a multigraph whose nodes are “tasks” and whose arcs are “streams”. In the denotation we choose, tasks and streams are mapped into atoms. The relations we have in our respective design model consist of

- the unary relations $Task(x)$ and $Stream(x)$, and
- the binary relation $Connect(x, y)$

which indicate that the directed arc representing stream x terminates at the node representing task y (or that the directed arc representing stream y begins at the node representing task x .)

It is easy to see that the general form of programs that parses Figure 4 is that of a directed multi-path, as depicted in Figure 5, and that the design models of these programs have the general form illustrated in formula (6).

$$\begin{aligned} &Connect(T_1, S_{1,1}), \dots Connect(T_1, S_{1,N1}), \dots \\ &Connect(S_{1,1}, T_2), \dots Connect(S_{1,N1}, T_2), \\ &\dots \\ &Connect(T_k, S_{k,1}), \dots Connect(T_k, S_{k,Nk}), \dots \\ &Connect(S_{k,1}, T_{k+1}), \dots Connect(S_{k,Nk}, T_{k+1}) \end{aligned} \quad (6)$$

Corollary 5. “Pipes and Filters” is *intensional* and *non-local*.

It is obvious that an unbounded number of programs satisfy formula (6); therefore, it is intensional. To show that (6) is non-local, observe that, for any design model that satisfies (6) we can add a new task T_{k+2} that is not connected to any other task by a pipe. Such addition violates the style’s specification (in particular, there would be no $Connect$ relation involving T_{k+2}). Since this is a violation *wherever* it occurs, the Pipes and Filters style must be non-local.

6.3 Law of Demeter

We have now shown that two classic architectural styles meet our criteria for being “architectural” as expected. But our criteria also turn up some less expected results: The *Law of Demeter* [20] was created as a *design heuristic*. It was introduced to simplify modifications to a program and to reduce its complexity. The informal description of the law for functions is given in Table 4.

We may formulate the language of Table 4 as follows:

$$\begin{aligned} &\forall f_1, f_2, c_1, c_2 \bullet \\ &Member(f_1, c_1) \wedge \\ &Member(f_2, c_2) \wedge \\ &Invoke(f_1, f_2) \Rightarrow \\ &Member(c_2, c_1) \vee ArgOf(f_1, c_2) \vee c_1 = c_2 \end{aligned} \quad (7)$$

Evidently, formula (7) has infinitely many instances, hence it is intensional. It is also non-local. To prove this, observe that any program that satisfies (7) can be expanded with source code that violates the Law, such as demonstrated by the C++ source code in Table 5.

Table 4. Law of Demeter for functions

For all classes C , and for all methods M attached to C , all objects to which M sends a message must be instances of classes associated with the following classes:

- The argument classes of M (including C).
- The instance variable classes of C .

(Objects created by M , or by functions or methods that M calls, and objects in global variables, are considered arguments of M .)

* In absence of an explicit name, we use the initials of the title of [6] with reference to the formalism.

Table 5. An add-on to a C++ program which violates the Law of Demeter

```

struct NewName1 {
    void foo();
};
struct NewName2 {
    NewName1 y;
};
class NewName3 {
    NewName2 x;
    void bar() {
        x.y.foo();
    }
};

```

In conclusion, the Law of Demeter, created as a design rule, can be better characterized as an *architectural* rule. The Law is not be limited to one part of the system but must be satisfied throughout. In practice, this means that any system using the Law of Demeter must create appropriate architectural practices to enforce it via coding standards, design walkthroughs, tool support, etc.

This example demonstrates the benefit of making our distinctions explicit and the power of rendering them precise, without which we would be unable to classify the Law of Demeter conclusively.

7. Analysis

Clearly, results obtained from the case studies in Sections 5 and 6 are not coincidental. The same line of reasoning used for the Factory Method can be used for many other (if not all) of the design patterns in [14], as well as for the architectural styles by Garlan and Shaw [16]. Examples drawn from other formal languages proposed for the specification of design patterns, such as *Constraint Diagrams* [19], *DisCo* [22], and *Contracts* [17], bring forth sample specifications, are clearly *intensional* as well. This motivates the following hypothesis:

The hypothesis of intensional specifications. All “design patterns” and “architectural styles” are intensional.

A direct proof of this requires a formalization of “all” design patterns and architectural styles. The first problem with this is that no given catalogue purports to contain “all” patterns and styles, nor do we expect such a catalogue to be possible (except perhaps in the analytic sense.) Another problem arises from the mostly informal definitions given to patterns and styles. Limited attempts have been made to prove this hypothesis (e.g., [13] [9]), but the proofs provided cannot cover *all* known patterns and styles. That is why the “hypothesis of intensional specifications” remains a hypothesis.

7.1 Specific “Designs” and “Architectures”

With the increase in popularity of the terms and their proliferation in the literature, they often appear in a concrete context, such as “the design of this program...” This usage implies that these terms can also be used with reference to extensional specifications, but only when referring to a concrete program. We suggest that “the design of program *x*” refers to the *instance* implemented in a program of a general design rule (e.g., design pattern). Since instances are extensions, this resolves the apparent difficulty in the intension/locality thesis.

7.2 UML Class Diagrams

Since UML is used widely as a design and architectural notation, it is of particular interest to understand the place of UML diagrams in the classification we introduced: Are they local? Intensional?

Despite the widespread attempts towards rendering the notation with well-defined semantics (e.g., the research group known as *pUML* [27]), most types of UML diagrams have no well-defined semantics. Thus, our discussion here is largely informal, assuming that any formal interpretation for class diagrams will be consistent with the informal semantics.

In terms of design models, we can assume that any such interpretation will associate class icons with atoms of type *class*, operations with atoms of type *method*, as well as provide us with a specification of a set of associated relations. It is easy to see why the specification given by a class diagram is local; but is it intensional?

To answer this question, note that a UML class diagram is commonly viewed as an under-specification, namely, an incomplete specification; the actual implementation of the diagram may have any number of additional elements that are not mentioned in the diagram. Under this assumption, UML class diagrams are intensional, since there exists an unbounded number of elements that can be added to any implementation.

Unlike the formulas used in our examples, however, the abstraction that class diagrams provide is of the most rudimentary type, that is, by omitting information but without using free variables. A UML diagram provides no information on the elements that are not explicitly described. Thus, class diagrams are intensional only in a trivial sense, in the same way that the code excerpts in Table 2 and Table 3, if taken not as complete programs but just as excerpts thereof, are intensional. Clearly, this sense is unlike the way architectural and design specifications are intensional. The following definition facilitates this distinction:

Definition IV. An intensional specification φ is **quasi-extensional** iff the set of design models that satisfy φ , has a single lower bound with respect to the partial-ordering relation “subsumption”.

It is trivial to show that subsumption induces partial ordering on a set of design models. Definition IV, however, assigns specific importance to sets of design models that contain design model such that all other members subsume it.

Corollary 6. UML class diagrams are quasi-extensional.

Figure 6 illustrates the proof to this corollary. It is trivial to show also that none of the intensional specifications we quoted above is quasi-extensional.

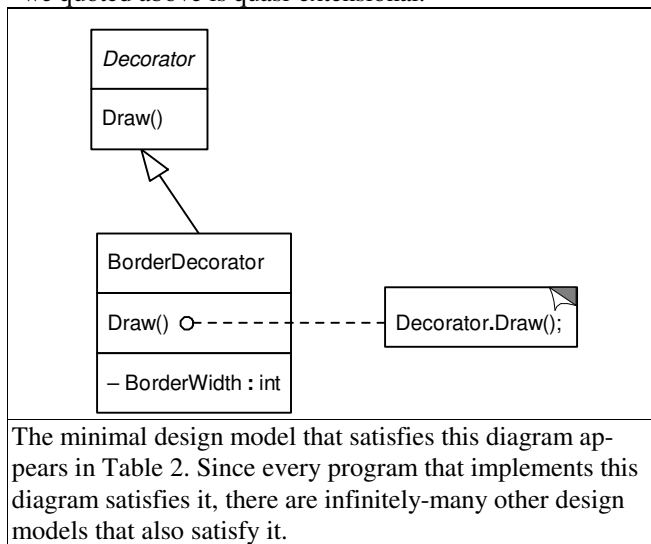


Figure 6. The design models that satisfy a UML class diagram

8. Conclusions

We have provided a sound formal basis for the distinction between the terms *architecture*, *design*, and *implementation*, called the intension/locality thesis, and established it upon best practices. What are the consequences of precisely knowing the differences between the terms architecture, design and implementation? Among others, these distinctions facilitate –

- determining what constitutes a uniform program, e.g., a collection of modules that satisfy the same architectural specifications;
- determining what information goes into *architecture* documents and what goes into *design* documents;
- determining what to examine and what to not examine in an architectural evaluation or a design walkthrough;

- understanding the distinction between local and non-local rules, i.e., between the design rules that need be enforced throughout a project versus those that are of a more limited domain.

For example, in the industrial practice of software architecture, many statements that are said to be “architectural” are in fact local, e.g., *both tasks A and B execute on the same node*, or *task A controls B*. Instead, a truly architectural statement would be, for instance, *for each tasks A,B which satisfy some property φ , A and B will execute on the same node and Control(A,B)*. More generally, for each specification we should be able to determine whether it is a *design* statement, describing a purely local phenomenon (and hence of secondary interest in documentation, discussion, or analysis), or whether it is an instance of an underlying, more general rule. (*)

Acknowledgements

Many thanks to Jens Jahnke and Alejandro Allievi for their comments. We thank Mary J. Anna for her inspiration. This research was supported in part by the Natural Sciences and Engineering Research Council, Canada and by the U.S. Department of Defense.

References

- [1] J. Barwise, ed., (1977). *Handbook of Mathematical Logic*. Amsterdam: North-Holland Publishing Co.
- [2] L. Bass, P. Clements, R. Kazman (1998). *Software Architecture in Practice*. Reading, MA: Addison Wesley Longman, Inc.
- [3] G. Booch, I. Jacobson, J. Rumbaugh (1999). *The Unified Modeling Language Reference Manual*. Reading, MA: Addison-Wesley.
- [4] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, M. Stal (1996). *Pattern-Oriented Software Architecture – A System of Patterns*. New York, NY: Wiley and Sons.
- [5] J. Coplien, D. Schmidt, eds. (1995). *Pattern Languages of Program Design*. Reading, MA: Addison-Wesley.
- [6] T. R. Dean, J. R. Cordy. "A Syntactic Theory of Software Architecture." *IEEE Trans. on Software Engineering* 21 (4), Apr. 1995, pp. 302–313.
- [7] F. DeRemer, H. H. Kron. "Programming-in-the-large versus programming-in-the-small." *IEEE Trans. in Software Engineering* 2 (2), June 1976, pp. 80–86.
- [8] A. H. Eden. "Formal Specification of Object-Oriented Design." *International Conference on Multidisciplinary Design in Engineering CSME-MDE 2001*, Nov. 21–22, 2001, Montreal, Canada.

* This final example was suggested by an anonymous ICSE reviewer.

- [9] A. H. Eden. "A Theory of Object-Oriented Design." *Information Systems Frontiers* 4 (4), Nov.—Dec. 2002. Kluwer Academic Publishers.
- [10] A. H. Eden. "LePUS: A Visual Formalism for Object-Oriented Architectures." *The 6th World Conference on Integrated Design and Process Technology*, Pasadena, CA, June 26—30, 2002.
- [11] A. H. Eden, Y. Hirshfeld. "Principles in Formal Specification of Object Oriented Architectures." *CASCON 2001*, Nov. 5—8, 2001, Toronto, Canada.
- [12] A. H. Eden, R. Kazman (2003). "On the Definitions of Architecture, Design, and Implementation". Technical report CSM-377, January 2003, Department of Computer Science, University of Essex, United Kingdom.
- [13] P. van Emde Boas (1997). "Resistance Is Futile; Formal Linguistic Observations on Design Patterns." Research Report no. CT-19997-03, The Institute for Logic, Language, and Computation, Universiteit van Amsterdam.
- [14] E. Gamma, R. Helm, R. Johnson, J. Vlissides (1994). *Design Patterns: Elements of Reusable Object Oriented Software*. Addison-Wesley.
- [15] D. Garlan, R. Monroe, D. Wile (1997). "ACME: An Architectural Description Interchange Language." *Proceedings of CASCON'97*. Toronto, Ontario.
- [16] D. Garlan, M. Shaw (1993). "An Introduction to Software Architecture." In V. Ambriola, G. Tortora, eds., *Advances in Software Engineering and Knowledge Engineering*, Vol. 2, pp. 1—39. New Jersey: World Scientific Publishing Company.
- [17] R. Helm, I. M. Holland, D. Gangopadhyay. "Contracts: Specifying Behavioral Compositions in Object-Oriented Systems." *Proceedings OPP-SLA/ECOOP*, Oct. 21—25, 1990, Ottawa, Canada.
- [18] R. Kazman. "A New Approach to Designing and Analyzing Object-Oriented Software Architecture." Guest talk, *Conference On Object-Oriented Programming Systems, Languages And Applications – OOPSLA*, Nov. 1—5, 1999, Denver, CO.
- [19] A. Lauder, S. Kent. "Precise Visual Specification of Design Patterns." *Proceedings of the 12th ECOOP, Brussels, Belgium*, July 1998. LNCS 1445. Berlin: Springer-Verlag.
- [20] K. Lieberherr, I. Holland, A. Riel (1988). "Object-oriented programming: an objective sense of style." *Conference proceedings OOPLA'88*, San Diego, CA, pp. 323—334.
- [21] D. C. Luckham et. al. "Specification and Analysis of System Architecture Using Rapide." *IEEE Trans. on Software Engineering* 21 (4), Apr. 1995, pp. 336—355.
- [22] T. Mikkonen (1998). "Formalizing Design Patterns." *Proceedings of the International Conference on Software Engineering*, April 19—25, 1998, pp. 115—124. Kyoto, Japan.
- [23] R. T. Monroe, A. Kompanek, R. Melton, D. Garlan. "Architectural Styles, Design Patterns, and Objects." *IEEE Software* 14(1), Jan. 1997, pp. 43—52.
- [24] D. E. Perry, A. L. Wolf (1992). "Foundation for the Study of Software Architecture." *ACM SIGSOFT Software Engineering Notes* 17 (4), pp. 40—52.
- [25] Popkin Software (2000). *System Architect 2001*. New York, NY: McGraw-Hill.
- [26] R. Prieto-Diaz, J. Neighbors. "Module Interconnection Languages." *Journal of Systems and Software* 6 (4), 1986, pp. 307—334.
- [27] *The Unambiguous UML Consortium* page: www.cs.york.ac.uk/puml/
- [28] T. Quatrani (1999). *Visual Modelling with Rational Rose 2000 and UML, Revised*. Reading, MA: Addison Wesley Longman, Inc.
- [29] D. C. Schmidt, M. Stal, H. Rohnert, F. Buschmann (2000). *Pattern-Oriented Software Architecture*, Vol. 2: Patterns for Concurrent and Networked Objects. New York, NY: John Wiley & Sons, Ltd.
- [30] SEI (2002). Carnegie Mellon's *Software Engineering Institute*. <http://www.sei.cmu.edu>.
- [31] M. Shaw, D. Garlan (1996). *Software Architecture: Perspectives on an Emerging Discipline*. Upper Saddle River, NJ: Prentice Hall.
- [32] J. M. Spivey (1989). *The Z Notation: A Reference Manual*. New Jersey: Prentice Hall.